

Modeling molecules with constraints

Charles F. F. Karney* and Jason E. Ferrara

Sarnoff Corporation, Princeton, NJ 08543-5300

(Dated: August 17, 2005)

Techniques for simulating molecules whose conformations satisfy constraints are presented. A method for selecting appropriate moves in Monte Carlo simulations is given. The resulting moves not only obey the constraints but also maintain detailed balance so that correct equilibrium averages are computed. In addition, techniques for optimizing the evaluation of implicit solvent terms are given.

Keywords: molecular simulation, constrained moves, energy evaluation

I. INTRODUCTION

When attempting to compute thermodynamic quantities with a molecular simulation, we are frequently confronted with the problem of sampling in a high-dimensional configuration space. The dimensionality of this space is given by the number of degrees of freedom for the molecular system. Techniques which lower the number of degrees of freedom will increase the efficiency of the thermodynamic sampling—provided, of course, that these techniques are physically justified. Thus, an implicit solvent model may be used to eliminate the degrees of freedom associated with the solvent molecules; the standard chemical force fields replace the electron charges with atom-centered partial charges thereby removing the electrons' degrees of freedom. Further reductions in dimensionality are possible by imposing constraints on the relative positions of the atoms in a molecule. Thus we might specify that the bond lengths and bond angles in a molecule are fixed and only the torsion angles are allowed to vary. It is such a scenario that we examine in this paper. We address two aspects of this problem: how to move a molecule subject to constraints in order to allow equilibrium averages to be computed using the canonical-ensemble Monte Carlo method [1] and how to evaluate the energy efficiently.

The imposition of constraints in molecular modeling has been extensively studied [2, §3.3.2, §15.1]. Let us start by elucidating the difference in the treatment of hard constraints in molecular dynamics and Monte Carlo simulations. We treat hard constraints by taking the limit where the “spring constant” for the hard degrees of freedom is infinite. Molecular dynamics simulations then consider the evolution of the resulting system over a finite time. On the other hand, if we wish to determine the equilibrium properties of a system using the Monte Carlo method, then we need to consider averaging over sufficiently long times to allow equipartition of energy among all the degrees of freedom of a system. This is, of

course, an example of nonuniform limits. We are interested in taking both $\tau_{\text{sim}} \rightarrow \infty$ and $\tau_{\text{equ}} \rightarrow \infty$, where τ_{sim} is the representative simulation time and τ_{equ} is the equipartition time (which is proportional to the “stiffness” of the constraints). In constrained molecular dynamics, we take the limit $\tau_{\text{equ}} \rightarrow \infty$ first, which prevents equipartition from occurring; whereas in equilibrium statistical mechanics, we take $\tau_{\text{sim}} \rightarrow \infty$ first and this allows energy equipartition. If we are attempting to compute an equilibrium quantity, such as the free energy of binding, it is essential to allow energy equipartition. Understanding this distinction explains the apparently contradictory results for constrained and unconstrained averages for a flexible trimer [2, §15.1].

One way of understanding the constrained equilibrium system is to consider how the equilibrium distribution varies as the constraint is imposed. In the limit, the distribution collapses to a lower-dimensional sub-manifold of configuration space. However, this sub-manifold has a “thickness” that depends on the details of the constraint term and, consequently, Monte Carlo moves for the constrained system need to reflect this thickness in order to sample the distribution correctly. As a consequence, we will need to specify the functional form of the constraint energy and the constraint is no longer a purely geometrical object. At first glance, this would appear to complicate further the already complex algebra of constrained motions [3]. However, we will propose an algorithm for making moves which is simple to implement and which automatically ensures that the correct equilibrium averages are computed.

The second half of the paper considers a mundane—but nevertheless important—problem, namely how to evaluate the energy of a molecule made up of rigid subcomponents. We propose a consistent framework for avoiding the computation of constant terms and for imposing energy cutoffs. We extend this to the computation of the generalized Born solvation term and we describe a simple method for computing the solvent accessible surface area which has a bounded error.

*Electronic address: ckarney@sarnoff.com

II. GENERALIZED MONTE CARLO MOVES

We begin by assembling some techniques for combining Monte Carlo moves. We define an “ E move” as an ergodic move which preserves $\exp(-\beta E)$ as the invariant distribution, where $\beta = 1/(kT)$ and k is the Boltzmann constant. (Here, “ergodic” implies that the move allows all relevant portions to configuration space to be explored.) Note that a zero move has a uniform invariant distribution. A typical zero move samples a new configuration from a distribution which satisfies the symmetry requirement $p(\Gamma'; \Gamma) = p(\Gamma; \Gamma')$, where $p(\Gamma'; \Gamma)$ is the probability density of picking a new configuration of Γ' given a starting configuration of Γ . Clearly a sequence of n E moves is itself an E move. From the central limit theorem, a sequence of n zero moves is equivalent, in the limit of large n , to selecting the new configuration from a multi-dimensional Gaussian.

Instead of carrying out the n moves with a given energy $E(\Gamma)$, we can consider the case where the energy is given by $E_\lambda(\Gamma)$ which depends continuously on the parameter λ . A sequence of n E_λ moves where λ is varied *adiabatically* in such a way that its initial and final values are λ_0 is an E_{λ_0} move. This follows because adiabatically varied systems are always in equilibrium with the instantaneous value of λ [4, §11]. Each E_λ move is carried out at a *fixed* λ and λ is varied between the moves. In order to satisfy the adiabatic condition, we will need to take n large.

A move from Γ to Γ' may be subjected to “Boltzmann acceptance with energy E ”. This involves accepting the move (Γ' is the new state) with probability $M(x)$ and otherwise rejecting the move (Γ is the new state). Here $M(x)$ is a function satisfying $0 < M(x) \leq 1$ and $M(x)/M(-x) = \exp(-x)$ with $x = \beta(E(\Gamma') - E(\Gamma))$. Usually we take $M(x) = \min(1, \exp(-x))$; however other choices, e.g., the Fermi function, $M(x) = 1/(1 + \exp(x))$, are possible [5].

Consider an E_1 move from Γ to Γ' followed by an Boltzmann acceptance using E_2 . This compound move is an $(E_1 + E_2)$ move. The proof follows as a special case of the “multiple time-step” (MTS) method [6, §II] or, alternatively, as a special case of early rejection [2, §14.3.2]. If the E_1 move was already a “rejected” move, i.e., $\Gamma' = \Gamma$, then the Boltzmann test involving E_2 automatically “succeeds” ($M(0) = 1$). Thus E_2 does not need to be evaluated in this case.

These results allow us to generalize the MTS method by splitting the energy into m terms (instead of just two),

$$E(\Gamma) = \sum_{l=1}^m E_l(\Gamma).$$

The method is defined recursively as follows: a level-0 move is defined to be a zero move; a level- l move, with $l > 0$, is defined to be n_{l-1} level- $(l-1)$ moves the result of which is subjected to Boltzmann acceptance using E_l . By induction, we see that a level- l move is an \mathcal{E}_l move, where

$$\mathcal{E}_l(\Gamma) = \sum_{\nu=1}^l E_\nu(\Gamma).$$

It follows that a level- m move is an \mathcal{E}_m move, i.e., an E move. Typically we sample the zero moves from a Gaussian and we take $n_0 = 1$. Standard Monte Carlo [1] is given by $m = 1$ and $n_1 = 1$. Standard MTS [6] is recovered with $m = 2$. The early rejection method [2, §14.3.2] is recovered with $n_l = 1$ (for all l). Note that a level- m move entails $\prod_{\nu=1}^{m-1} n_\nu$ level- l moves. At any stage in the recursion, we have the freedom to vary some of the components of $E(\Gamma)$ adiabatically.

In the following sections, we apply these techniques to constrained molecules. In simple cases, we can apply the MTS method semi-analytically to derive a correct constrained move. In more complicated cases, we apply the adiabatic technique to lift and to reapply the constraint.

III. STIFF MOLECULES

A constrained molecule is a mathematical idealization of a real system in which some degrees of freedom are stiff, i.e., the associated energies are large. Thus we can split the energy into “hard” (h) and “soft” (s) components,

$$E(\Gamma) = E_h(\Gamma) + E_s(\Gamma),$$

where Γ is the configuration of the system. For example, let us assume that an all-atom force field, such as Amber [7], provides an accurate description of the system. (We recognize, of course, that present-day force fields are only approximate. However, our purpose here is to make the connection between an all-atom representation and a simpler rigid representation and, in this context, the details of the all-atom model are of secondary importance.) Then E_h might represent the bond stretching and bond bending terms, while E_s is given by the other terms (bond torsion and the non-bonded energies).

The constrained limit is now given by $E_h \rightarrow \infty$. Before we consider this limit, it is useful to examine how the stiff system may be treated. Conventional Monte Carlo is inefficient because, in order to have a reasonably large acceptance rate, the step-size needs to be set to a small value (determined by E_h) so that diffusion in the soft directions is very slow. However, we can apply MTS Monte Carlo in this case with $E_1 = E_h$ and $E_2 = E_s$.

Let us apply this method to a system of “rigid” molecules, e.g., water molecules, taking E_h to include the intra-molecular energies (responsible for maintaining the rigidity) and E_s to include the inter-molecular energies. Suppose the level-0 moves consist of symmetrically displacing the atoms in each molecule. The result of the n_1 level-1 Monte Carlo steps will clearly be a symmetric, independent, and nearly rigid displacement (translation and orientation) of each molecule. This configuration is then subjected to Boltzmann acceptance with the inter-molecular energies. In this case, we can easily pass to the constrained limit (with exact rigidity), merely by ensuring that the trial (level-1) moves of the molecules are rigid. In this case, we have just rederived the “standard” move for a system of rigid molecules.

In order to illustrate the application to flexible molecules, we shall treat the molecules as being made up of several rigid

subunits or “fragments” connected by flexible bonds. However we are interested in the limit where the inter-fragment bonds constrain the relative motions of fragments in certain ways, either by fixing the bond lengths (allowing the bond angles and bond dihedrals to vary) or by fixing the bond lengths and bond angles (allowing the bond dihedrals to vary). Such a model is adequate to describe a wide range of interesting organic molecules including proteins and drug-like ligands. We assume that the rigidity of the fragments is imposed only by intra-fragment energy. If other terms (e.g., an improper torsion term involving atoms from two fragments) contribute to the rigidity of a fragment, then we shall treat such terms as additional inter-fragment energies.

We apply the generalized MTS method to this system with $m = 3$, the intra-fragment energy given by E_1 , the inter-fragment bond constraints given by E_2 , and with E_3 accounting for all the other energies. The argument given above allows us to pass to the limit of strictly rigid fragments. The method is then equivalent to a standard MTS method where the “elementary” moves consist of rigid displacements of each fragment which are Boltzmann accepted with energy $E_h = E_2$. A sequence of $n = n_2$ such moves are made with the result Boltzmann accepted with energy $E_s = E_3$. A possible prescription [8, §VII] for the rigid displacements of the fragments is to translate the fragment by a vector sampled from an isotropic 3-dimensional Gaussian and to rotate the fragment by $|\mathbf{s}|$ about an axis $\hat{\mathbf{s}}$ where \mathbf{s} is a “rotation vector” also sampled from an isotropic 3-dimensional Gaussian. The variances for the two Gaussians should be adjusted so that the translational and rotational components result in comparable displacements of the atoms of the fragment.

Provided that the inter-fragment constraint terms E_h are sufficiently stiff, it is not important to include a detailed model of these terms; because the motion will take place near the bottom of the constraint potential well, a harmonic (i.e., quadratic) approximation to the constraint potential will suffice. On the other hand, if the stiffness of the constraint energy depends on any of the soft degrees of freedom, it is important that this effect be included.

It is frequently the case that E_h may be computed much more rapidly than E_s . For example, when imposing bond constraints on a molecule, E_h requires $O(N)$ computations, where N is the number of atoms, while E_s requires $O(N^2)$ computations for the electrostatic and implicit solvation energies. Thus we might be able to take n reasonably large and still have the computational cost dominated by the evaluation of $E_s(\Gamma)$.

In order to realize the full benefits of imposing constraints we need to pass to the constrained limit ($E_h \rightarrow \infty$). In this limit, the motion collapses onto a lower-dimensional submanifold in configuration space. Unfortunately, in contrast to the case of rigid molecules, we cannot appeal to symmetry to enable us to take this limit analytically. Instead, we use the adiabatic technique.

IV. ADIABATICALLY VARYING THE STIFFNESS

Let us rewrite the energy of the system, multiplying the $E_h(\Gamma)$ by T/T^* , where T is the temperature of the system, and T^* is a “constraint” temperature. The Boltzmann factor $\exp(-\beta E)$, will then have the form

$$\exp(-\beta E) = \exp(-\beta E_s - \beta^* E_h)$$

where $\beta^* = 1/(kT^*)$.

In our application, where we are interested in the constrained limit $T^* \rightarrow 0$, a direct application of the MTS method leaves us with two bad choices. If we take T^* to be sufficiently small that we can consider the constraints to be satisfied, we will have to choose the step size for the E_h moves to be so small that the change in configuration after n E_h moves will be small. On the other hand, letting T^* be sufficiently large to allow moves will result in configurations where the constraints are poorly satisfied.

We overcome this difficulty by regarding T^* as a parameter (taking the place of λ) and by adiabatically varying T^* from zero (where the constraints are satisfied but MTS is ineffective at making moves) to a finite value (where the constraints are relaxed and MTS becomes effective) and back to zero again (to reimpose the constraints). During the course of changing T^* , we make n E_h moves (each with the instantaneous value of T^*). The effect of these n moves will be an E_h move with $T^* = 0$, i.e., a move which satisfies the E_h constraint.

It remains to give a recipe for varying T^* . As we vary T^* , we would naturally adjust the step size for the moves in such a way that the number of steps needed to equilibrate the system is a constant, suggesting that we vary T^* exponentially. We therefore pick

$$T_i^* = \begin{cases} T_A^* \exp(\alpha(i-1)), & \text{for } 0 < i \leq m, \\ T_A^* \exp(\alpha(n-i)), & \text{for } m < i \leq n, \end{cases}$$

where we have taken $n = 2m + 1$ and where T_i^* is the constraint temperature used for the i th E_h move, $T_0^* = T_A^*$ is some temperature sufficiently small that we can consider the constraints to be exactly satisfied, and α is the rate of increase of the temperature which should be sufficiently small that the adiabatic condition is satisfied. Even though T_A^* and α are small, we can pick n sufficiently large that $T_{m+1}^* = T_B^* = T_A^* \exp(\alpha m)$ is finite.

In addition, we choose the step size for the i th E_h move to be $d_i = k\sqrt{T_i^*}$ where k is a constant. In traditional Monte Carlo, we normally pick k to maximize the diffusion rate which at the i th step is roughly

$$D_i = \frac{\langle (\Gamma_i - \Gamma_{i-1})^2 \rangle}{2} \sim \frac{1}{2} A d_i^2,$$

where A is the mean acceptance rate and $\langle \dots \rangle$ denotes an ensemble average. Maximizing the diffusion rate usually results in a rather small acceptance rate $A \sim 0.1$ because rare large steps can lead to faster diffusion than frequent small steps. However, in our application, where we want the system to remain in equilibrium as we vary the temperature, rare large

steps are *bad*. So we pick k to maximize AD_i and this will usually result in $A \sim 0.5$. Note that for a given k , we have

$$D_i \sim CT_i^*,$$

where C is constant provided that the step size is not too large. The overall diffusion can be estimated by summing over the n steps,

$$D = \frac{\langle (\Gamma_n - \Gamma_0)^2 \rangle}{2} = \sum_{i=1}^n D_i \sim 2CT_B^*/\alpha,$$

where we have assumed that successive steps are uncorrelated and we have taken $\alpha \ll 1$ and $T_B^* \gg T_A^*$. We should select parameters, α and T_B^* , in order to adjust D so that the E_s acceptance rate is $O(1)$.

This method includes internal diagnostics to verify that α is small enough. We define $\overline{\dots\uparrow}$ (resp. $\overline{\dots\downarrow}$) as the average of a quantity over the steps where T_i^* is increasing, i.e., $i \leq m+1$ (resp. decreasing, i.e., $i > m+1$). We monitor $\overline{E_h(\Gamma_i)/T_i^*\uparrow}$ and $\overline{E_h(\Gamma_i)/T_i^*\downarrow}$ and demand that both should be close to the equilibrium value of $N/2$ (where N is the number of hard degrees of freedom). If α is too large, then we would find

$$\begin{aligned} \overline{E_h(\Gamma_i)/T_i^*\uparrow} &\ll N/2, \\ \overline{E_h(\Gamma_i)/T_i^*\downarrow} &\gg N/2. \end{aligned}$$

In particular, if the final $E_h(\Gamma_n)$ is many times T_A^* , then the configuration is “hung up” and does not obey the constraints. If this happens frequently, the simulation needs to be rerun with a smaller setting for α ; if, on the other hand, it happens only rarely, we would merely reject the step. We can also monitor the mean acceptance rates $\overline{A\uparrow}$ and $\overline{A\downarrow}$. These should be about the same; however, if α is too large, we will find $\overline{A\uparrow} \gg \overline{A\downarrow}$.

A useful guideline for picking T_A^* is that once the n E_h moves are completed and the system is presumably equilibrated to T_A^* , we should be able to enforce the constraints by setting $T^* = 0$ (using any convenient energy minimization technique) with a negligible change in the configuration, e.g., with a negligible change in $E_s(\Gamma)$.

V. PAIRWISE TERMS IN ENERGY

Having made an adiabatic move using E_h , the final step is to accept the move depending on the change in E_s . We wish to compute this energy as efficiently as possible by using the rigidity of the fragments. Force fields such as Amber [7] include two types of energies: interactions between atoms (the electrostatic and Lennard-Jones terms) and bond energies (stretch, bend, and torsion). Since the number of terms in non-bonded energies typically scales as $O(N^2)$ where N is the total number of atoms in the system, while the number of bond terms scales as $O(N)$, we concentrate on optimizing the evaluation of the non-bonded terms. In our case where the molecules consist of rigid fragments connected by flexible bonds we need only include the bond terms contributed

by the much smaller number of inter-fragment bonds. Furthermore, we need only include the energy contributed by the “free” components of such bonds. Thus, if the lengths and angles of such bonds are constrained, then we need only include the torsion energy in $E_s(\Gamma)$.

We start by assuming that the non-bonded energy terms can be expressed as a sum over atom pairs. This applies to the electrostatic and Lennard-Jones terms in Amber [7]. However, implicit solvent models have a more complex structure and we consider these in the next section.

Suppose our molecular system consists of N atoms. These atoms are grouped into M molecules and we denote M_l as the set of atoms making up the l th molecule. Similarly, the atoms are divided into F rigid fragments and we denote F_a as the set of atoms making up the a th fragment. A typical pairwise energy term can then be written as

$$E_g(\Gamma) = \sum_{0 < i < j \leq N} C_{g,ij} f_g(r_{ij}),$$

where g denotes the type of energy term (electrostatic or Lennard-Jones), i and j are atom indices, r_{ij} is the distance between atoms i and j , f_g is some function of distance, and $C_{g,ij}$ is a coefficient which depends on the atoms but not on their positions. Thus for electrostatic interactions, $C_{g,ij}$ depends on the partial charges on the two atoms (assumed to be constant in Amber) and on the bonding relation between the atoms. Physical energy functions satisfy $\lim_{r \rightarrow \infty} f_g(r) = 0$. When the fragments are separated sufficiently, we have

$$E_g \rightarrow E_{g0} = \sum_{0 < a < b \leq F} \sum_{\substack{i < j \\ i,j \in F_a}} C_{g,ij} f_g(r_{ij}),$$

which is independent of Γ . It is convenient to choose E_{g0} as the “origin” for the E_g , i.e., we compute only

$$E_{g1} = E_g - E_{g0} = \sum_{0 < a < b \leq F} \sum_{\substack{i \in F_a \\ j \in F_b}} C_{g,ij} f_g(r_{ij}).$$

We note that only energy differences enter into the computation of observable quantities, and so we are free to select the arbitrary origin for energies.

Let us consider the application of a small molecule (N_l atoms) interacting with a protein ($N_p \gg N_l$ atoms) where only some of the protein side chains near the binding site are allowed to move. By avoiding computing the interaction energy between atoms in the immobile portion of the protein, the above prescription reduces the computational cost from $O(N_p^2)$ to $O(N_l N_p)$.

This cost may still be too large and we can substantially reduce the cost by implementing energy cutoffs for the interactions. This is easily accomplished by multiplying $f_g(r_{ij})$ by a cutoff function, $c_g(r_{ij})$. A possible form for this cutoff function is

$$c_g(r) = \begin{cases} 1, & \text{for } r < r_{g1}, \\ 0, & \text{for } r \geq r_{g2}, \\ c_g(r_{g1}) \frac{r_{g2} - r}{r_{g2} - r_{g1}}, & \text{otherwise,} \end{cases}$$

with $r_{g1} \leq r_{g2}$, which linearly tapers the energy to zero over $[r_{g1}, r_{g2}]$. Other tapering functions can be employed, or, by choosing $r_{g2} = r_{g1}$, we can implement a sharp cutoff. This type of cutoff function implements a per-atom cutoff and is appropriate for energy terms which are additive at large distances, such as the Lennard-Jones potential. The electrostatic potential, however, involves substantial cancellation at large distances—two neutral molecules interact via a dipole-dipole term which varies as $1/r^3$, while the individual atom-atom terms decay as $1/r_{ij}$. In this case, we need to identify groups of atoms which should interact together. The residues of a protein provide a convenient grouping and we would typically assign all the atoms in a small-molecule ligand to a single group. Compatible with the usage for a protein, we refer to these groups as residues. For each residue, s , we define a center position, \mathbf{b}_s , most conveniently defined as the center of mass, and a radius, defined as the radius h_s of the sphere centered at \mathbf{b}_s which includes the van-der-Waals spheres of radius ρ_i of all the constituent atoms. We then apply a “per-residue” cutoff function multiplying the contribution from the residue pair (s, t) by $c_g(|\mathbf{b}_s - \mathbf{b}_t| - (h_s + h_t))$.

The values used for the cutoff radii, r_{g1} and r_{g2} , need to be evaluated based on the accuracy desired for the simulation. This can be determined by numerically determining the difference in the results (either for the energies directly or for some derived quantity such as binding affinity) between the finite- and infinite-cutoff energies. In applications to Monte Carlo codes, it is possible to carry out the sampling at an energy approximating the actual energy and to compensate for this when performing the canonical averages (which might be carried out on a subset of the Markov chain). In this case, the sampling energy might entail using shorter cutoffs than would be warranted on the basis of accuracy. Having determined suitable cutoffs, it is a simple matter to evaluate the energy avoiding treating atom pairs beyond the respective cutoffs. In the following, we treat electrostatic (e) interactions, with a per-residue cutoff, and Lennard-Jones (l) interactions, with a per-atom cutoff; furthermore we assume that $r_{e2} \geq r_{l2}$, i.e., the electrostatic interactions are longer range than the Lennard-Jones.

We first loop over all the atoms in each residue computing \mathbf{b}_s and h_s for all residues s . We then loop over all pairs of residues, $s \leq t$, skipping any pair whose atoms all belong to the same fragment or those for which $|\mathbf{b}_s - \mathbf{b}_t| \geq r_{e2} + h_s + h_t$. If the residue pair survives these tests, then all atom pairs (i, j) from different fragments are considered; if $s = t$, we restrict the pairs to $i < j$. All such pairs contribute to the electrostatic energy while those which satisfy $r_{ij} < r_{l2}$ contribute to the Lennard-Jones energy. There obviously is scope for additional optimization here. For example, the inner atom loop can be skipped if the second residue belongs to a single fragment which matches the fragment of a particular atom in the first residue.

Because of the way in which the cutoffs are applied, the result for the energy is independent of the assignment of atoms to residues for energy terms which use a per-atom cutoff. In addition, *differences* in the non-bonded energies are independent of the assignment of atoms to fragments. The energies

for assemblies of 3 or more molecules can be expressed in terms of the energies of 1 or 2 molecules. These provide useful checks on the implementation.

In some contexts it is useful also to define a “steric” energy term which is infinite if any atoms overlap (with some definition of a “hard” atom radius) and is zero otherwise. This provides a rapid check of new configurations—particularly when trying to “insert” a molecule during a grand canonical simulation [9] or when switching systems using the wormhole method [10]. A conservative definition of the hard atom radius is $0.55\rho_i$ for non-bonded atom pairs and $0.45\rho_i$ for 1-4 atom pairs. We skip the check for 1-2 and 1-3 pairs and for those atoms with $\rho_i = 0$. This energy term can be implemented in essentially the same way as described above but with scope for additional speedups. The cutoff radius in the residue-residue distance check can be replaced by 0. An additional atom-residue distance check can be used to avoid executing the inner atom loop if the outer atom is outside the sphere for the second residue. Finally, as soon as an overlap of hard spheres is detected the routine can immediately return an infinite result.

VI. IMPLICIT SOLVENT MODELS

We now turn to the computation of the energy term for implicit solvent models. We focus here on the generalized Born solvent models [11] and we have considered various implementations [12; 13; 14; 15; 16; 17]. Evaluating the solvation energy for a system of molecules with such models is typically orders of magnitude slower than computing the energy of the molecules in vacuum. The computation time is frequently compared to the time to compute the energy with an explicit solvent model (including $O(10^3)$ solvent molecules). However, such comparisons are misleading because implicit solvent models do not attempt to compute the energy of a particular configuration of solvent molecules but to compute the *free energy* of solvation, i.e., to average over all possible solvent configurations for a given configurations of solute molecules. Thus the chief benefit of an implicit solvent model is to reduce dramatically the number of degrees of freedom in the problem. In the generalized Born solvent models, the energy is written as the sum of two terms: a polar term which is usually called the “GB” term and a cavity term which is proportional to the solvent accessible surface area, the “SA” term.

The GB term involves long-range interactions and is the most costly to compute. We address the calculation of this term first. The basic expression is [11]

$$G_{\text{pol}} = -\frac{1}{2} \frac{1}{4\pi\epsilon_0} \left(1 - \frac{\epsilon_0}{\epsilon_s}\right) \sum_{i,j} q_i q_j f(r_{ij}, \alpha_i, \alpha_j), \quad (1)$$

where ϵ_s is the permittivity of the solvent, $f(r_{ij}, \alpha_i, \alpha_j) = [r_{ij}^2 + \alpha_i \alpha_j \exp(-r_{ij}^2/(4\alpha_i \alpha_j))]^{-1/2}$, and the double sum runs over *all* pairs of atoms (including $i = j$ and $i \leq j$). In eq. (1), α_i is the “generalized” Born radius of the i th atom, which is larger than the “bare” Born radius to account for the

fact that atoms close to i partially shield it from the solvent. G_{pol} represents the electrostatic energy required to solvate a pre-assembled group of molecules and thus this term is added to the vacuum electrostatic energy. The various implementations for the GB term differ in how α_i is computed.

For illustrative purposes, let us consider the model of Hawkins *et al.* [13; 14; 15; 16]. (With minor modifications, the technique is applicable to other GB models.) We express α_i as [13, eq. (10)]

$$\frac{1}{\alpha_i} = \frac{1}{\rho_i} - \sum_{j \neq i} \Delta_{ij}, \quad (2)$$

where ρ_i is the radius of atom i ,

$$\Delta_{ij} = \int_{\rho_i}^{\infty} \frac{dr}{r^2} H_{ij}(r; r_{ij}, \rho_j) \quad (3)$$

is the reduction in the effective inverse Born radius of atom i due to atom j . Here H_{ij} is the fraction of the area of a sphere of radius r centered on the i th atom eclipsed by a j th atom and is given by [13, eq. (12)]

$$H_{ij} = \begin{cases} \frac{\rho_j^2 - (r_{ij} - r)^2}{4r_{ij}r}, & \text{for } |r_{ij} - \rho_j| \leq r \leq r_{ij} + \rho_j, \\ 1, & \text{for } r < \rho_j - r_{ij}, \\ 0, & \text{otherwise } (r \geq r_{ij} \pm \rho_j). \end{cases}$$

Evaluating the integral in eq. (3) then gives

$$\Delta_{ij} = \begin{cases} 0, & \text{for } \rho_i > \rho_j + r_{ij}, \\ \frac{l_{ij} - u_{ij}}{2} - \frac{(r_{ij}^2 - \rho_j^2)(l_{ij}^2 - u_{ij}^2)}{8r_{ij}} \\ \quad - \frac{\ln(l_{ij}/u_{ij})}{4r_{ij}} + l'_{ij}, & \text{otherwise,} \end{cases}$$

where $u_{ij} = 1/(r_{ij} + \rho_j)$, $l_{ij} = 1/\max(\rho_i, |r_{ij} - \rho_j|)$, and $l'_{ij} = 1/\rho_i - 1/\max(\rho_i, \rho_j - r_{ij})$. The term l'_{ij} is only non-zero for $\rho_j > \rho_i + r_{ij}$, which is a possibility not considered in [13].

Clearly G_{pol} is no longer the sum of pairwise atom-atom contributions because the interaction of two atoms is affected by the modification of the dielectric environment by a third atom. However G_{pol} may be evaluated by two pair-wise operations carried out in sequence. The first evaluates the generalized Born radii α_i and the second computes the resulting electrostatic energy.

As with the treatment of the electrostatic and Lennard-Jones terms, we can seek to limit the computational cost of evaluating G_{pol} by the use of cutoff functions. Because eq. (1) provides the dielectric screening for the vacuum electrostatic term, it is important that the cutoff function multiplying $f(r_{ij}, \alpha_i, \alpha_j)$ exactly match that used for the electrostatic term.

We also introduce a cutoff in eq. (2) by multiplying Δ_{ij} by $c_b(r_{ij})$. A per-atom cutoff is justified since all the Δ_{ij} are positive. Because Δ_{ij} scales as r_{ij}^{-4} for large r_{ij} , the error introduced by $c_b(r_{ij})$ scales relatively slowly as r_{b1}^{-1} . In practice, this means we need to make r_{b1} reasonably large which

in turn means that the cost of evaluating G_{pol} in the case of a small ligand interacting with a protein is much larger than the cost for the electrostatic potential. In particular, the screening of the ligand may modify the Born radii of a large number of protein atoms and this unavoidably leads to a large number of pair contributions to eq. (1).

The procedure for computing the energy outlined in the previous section can now be modified to deal with the evaluation of G_{pol} . As before our “zero” energy is given by separating all the fragments of all the molecules infinitely far apart. We set up the calculation of a system of molecules by pre-computing α_{i0} which is given by eq. (2) with the sum restricting to include only the intra-fragment contributions (i.e., index j ranges only over atoms within the same fragment as atom i). We compute Δ_{ij} and Δ_{ji} together because they involve many of the same terms, allowing the loops to be restricted to $i < j$, and we apply the Born cutoff to the calculation of α_{i0} .

When computing the energy of a molecular system, we compute all the updates to the Born radii due to atoms in different fragments within the Born cutoff, applying the same techniques of lumping the atoms into residues described above (which allows the cutoff criteria to be applied to groups of atoms) and of restricting the loops to $s \leq t$ and, for $s = t$, to $i < j$. During this phase we mark all the residues which contain atoms with $\alpha_i \neq \alpha_{i0}$. We then make a second pass over the atoms to evaluate the terms in eq. (1). We use the $i \rightleftharpoons j$ symmetry of the summand to make the restrictions $s \leq t$ and, for $s = t$, $i \leq j$. In the innermost loop, we accumulate $q_i q_j f(r_{ij}, \alpha_i, \alpha_j)$ if i and j belong to different fragments. Otherwise, we add $q_i q_j [f(r_{ij}, \alpha_i, \alpha_j) - f(r_{ij}, \alpha_{i0}, \alpha_{j0})]$ and we can skip this evaluation if both $\alpha_i = \alpha_{i0}$ and $\alpha_j = \alpha_{j0}$. In addition, we can skip pairs of residues if all the atoms in each residue belong to the same fragment and if neither residue is marked as having modified Born radii.

Salt effects [18] are easy to include within this framework. A minor complication occurs in the GB model of Qiu *et al.* [12] because α_{i0} depends on the “volume” of the atoms and in this model the volume depends on the 1-2 bonded atoms which may belong to a different fragment. We account for this by assuming the presence of such bonded atoms with an ideal bond length. This is, therefore, only exact if the inter-fragment bonds are at their ideal lengths. Our treatment here may be considered as a generalization of the frozen atom approximation for GB/SA [19]. However, in our application we make all the approximations in the energy function and the resulting energy is then a “state variable” and simulations based on this are well behaved. In contrast the implementation of frozen atom approximation defines the energy so that it depends on the history of the system which may cause the simulation to exhibit unphysical properties.

VII. SOLVENT ACCESSIBLE SURFACE AREA

The other important contribution to the solvation free energy is the cavity term. This is obtained by placing spheres centered at each atom with radius $a_i = \rho_i + r_w$ where r_w is a nominal water radius (typically $r_w = 0.14$ nm). The cavity

term is given by

$$G_{\text{cav}} = \sum_i \sigma_i A_i,$$

where A_i is the “solvent accessible surface area” for the i th atom, i.e., the exposed surface area of the spheres around i which is not occluded by any other spheres and σ_i is the surface tension for the i th atom. (Typically σ_i is taken to be a constant independent of atom, $\sigma_i \approx 3 \text{ kJ mol}^{-1} \text{ nm}^{-2}$; however the method we describe does not require this assumption.) As before, the zero energy state is obtained by separating the fragments infinitely. The energy is then given by the additional occlusion of the surface that occurs as the fragments are assembled into molecules and the molecules brought into contact with one another.

The exact evaluation of this term is quite complex and for this reason a simple pairwise approximation has been developed [20]. However, the errors in this method are poorly quantified. This together with the fact that this term is typically small compared to the electrostatic terms in the energy lead us to develop a simple zeroth-order quadrature method. We select an accuracy level for the cavity calculation δ , e.g., $\delta = 0.1 \text{ kJ/mol}$. We prepare for the calculation of the cavity term by placing each fragment in a “template” position and we arrange a set of points on a sphere of radius a_i around each atom i . The number of points is chosen to be $N_i = \lceil 4\pi a_i^2 \sigma_i / \delta \rceil$. The points are distributed approximately uniformly around each sphere and the entire surface energy of the sphere, $4\pi a_i^2 \sigma_i$ is divided among the N_i points. (We will discuss the details of how to select the points and assign the energy later.) We next perform the intra-fragment occlusion by deleting all the points of atom i which are within a_j of some atom $j \neq i$. In this way each fragment is surrounded by a cloud of surface points each representing about δ of cavity energy.

In order to compute the cavity term for a particular molecular configuration we transform the surface points for each fragment from their template positions to their actual positions and make a copy of the cavity energies for each point. We consider all pairs of atoms (i, j) such that i and j are in different fragments and $r_{ij} < a_i + a_j$. We subtract from G_{cav} the energies of all the points on atom i that are within a_j of atom i and we set the energies of these points to zero (to avoid their being counted multiple times). The optimizations described above can be used: the application of a residue-residue cutoff (excluding residue pairs (s, t) with $|\mathbf{b}_s - \mathbf{b}_t| \geq 2r_w + h_s + h_t$), an atom-residue cutoff, and the treatment of the (i, j) and (j, i) terms together.

In practice, the cost of evaluating this term is small for $\delta \approx 0.1 \text{ kJ/mol}$. The error is proportional to δ and it is easy to benchmark a particular calculation by repeating it with smaller δ . The resulting G_{cav} is obviously a discontinuous function of configuration, jumping by $\pm\delta$ as points move in and out of the water spheres of other atoms. Thus it’s an inappropriate model for a molecular dynamics simulation. However, it yields satisfactory results for Monte Carlo simulations.

Let us return to the question of how to position the points on the atom sphere and how to divide the energy between

these points. Ideally, we would divide the energy of the sphere based on the area of Voronoi polygons around each point. The error will then be proportional to the maximum radius of the Voronoi polygons and the ideal distribution of points is the one which minimizes this maximum radius. This is the so-called “covering problem” for the sphere, i.e., how to cover a sphere with identical discs [21]. Unfortunately, there are no general solutions to this problem. So instead we divide the sphere into equal intervals of latitude and we divide each latitudinal interval longitudinally into approximately square regions. A point is placed at the center of each region and the area of the region is assigned to that point. Within each fragment, we alter the position of the pole from one atom to the next, in order to avoid the occlusion of many points simultaneously as fragments move relative to one another.

VIII. DISCUSSION

We have shown how to make Monte Carlo moves for a molecular system with constraints. Constraints are imposed in a realistic way ensuring that we obtain the right distribution corresponding to a thermodynamic equilibrium. We will still need to know this constrained distribution if we wish to make wormhole moves [10], because, in order to satisfy detailed balance, we require knowledge of the wormhole volumes and these include a factor proportional to the “thickness” of the constraint manifold. The adiabatic move involves, naturally, many evaluations of the constraint energy raising a concern that the implementation will be slow. In reality, the cost of evaluating the constraint energy is minuscule, particularly in comparison with the solvation energy, so it is possible to evaluate the constraint energy many thousands of times in the course of an adiabatic move with minimal impact on the overall running time. The method avoids much of the algebra associated with other ways of imposing constraints [22] and thus is more flexible and is easier to implement.

In the simple case of a molecule in which only a number of dihedral angles are allowed to vary, the movement of all the atoms in the molecule is bounded and thus the soft-energy acceptance probability is reasonably large. In contrast, the method where the dihedral angles are perturbed may lead, due to a lever effect, to large motions if the molecule itself is large.

This method can easily be generalized to do localized movements. Thus, we can tailor the random displacements of a protein to explore the movement of a single loop. Detailed balance is ensured if the random displacement is a function of the atom but not of its position. (The general case can be accommodated by a suitable factor in the acceptance probability.) This method of localized movements is more widely applicable than techniques such as “concerted rotations” [22; 23; 24]. Artificially fixing the positions of some atoms would, of course, mean that the moves would not be ergodic. This would be justified if we were interested in examining the restricted system and we would then require ergodicity over the restricted configuration space.

We have also considered how to optimize the evaluation of the energy in a system of molecules made up of rigid

fragments bonded together. This allows the use of implicit solvent at an acceptable cost. If the system is further constrained to allow only the variation of the torsion angle of the inter-fragment bonds (fixing the bond lengths and bond angles), then we should also consider modifying the force field to “loosen” the torsion energies to counteract the effect of the hard constraints on the other bond terms. Gō and Scheraga [25] show the importance of considering such an effect and Katrich *et al.* [26] have offered a prescription for converting a general force field to include this effect. Alternatively, we might consider re-parameterizing the torsion terms by carrying out constrained geometry optimizations of model molecules where the energy of the molecule is minimized with the dihedral angles fixed [27].

Acknowledgment

This work was supported by the U.S. Army Medical Research and Materiel Command under Contract No. DAMD17-03-C-0082. The views, opinions, and findings contained in this report are those of the author and should not be construed as an official Department of the Army position, policy, or decision. No animal testing was conducted and no recombinant DNA was used.

References

- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of state calculations by fast computing machines*, J. Chem. Phys. **21**, 1087–1092 (1953).
- [2] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, 2002), 2nd edition.
- [3] M. Fixman, *Classical statistical mechanics of constraints: A theorem and application to polymers*, Proc. Nat. Acad. Sci. **71**, 305–3053 (1974).
- [4] L. D. Landau and E. M. Lifshitz, *Statistical Physics*, vol. 5 of *Course of Theoretical Physics* (Pergamon Press, 1969), 2nd edition.
- [5] C. H. Bennett, *Efficient estimation of free energy differences from Monte Carlo data*, J. Comput. Phys. **22**, 245–268 (1976).
- [6] B. Hetényi, K. Bernacki, and B. J. Berne, *Multiple “time step” Monte Carlo*, J. Chem. Phys. **117**, 8203–8207 (2002).
- [7] W. D. Cornell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz, Jr., D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell, and P. A. Kollman, *A second generation force field for the simulation of proteins, nucleic acids and organic molecules*, J. Am. Chem. Soc. **117**, 5179–5197 (1995).
- [8] C. F. F. Karney, *Quaternions in molecular modeling*, Technical report, Sarnoff Corp. (June 2005). E-print arXiv:physics/0506177.
- [9] D. J. Adams, *Grand canonical ensemble Monte Carlo for a Lennard-Jones fluid*, Mol. Phys. **29**, 307–311 (1975).
- [10] C. F. F. Karney, J. E. Ferrara, and S. Brunner, *Method for computing protein binding affinity*, J. Comput. Chem. **26**, 243–251 (2005). E-print arXiv:cond-mat/0401348.
- [11] W. C. Still, A. Tempczyk, R. C. Hawley, and T. Hendrickson, *Semianalytical treatment of solvation for molecular mechanics and dynamics*, J. Am. Chem. Soc. **112**, 6127–6129 (1990).
- [12] D. Qiu, P. S. Shenkin, F. P. Hollinger, and W. C. Still, *The GB/SA continuum model for solvation: A fast analytical method for the calculation of approximate Born radii*, J. Phys. Chem. A **101**, 3005–3014 (1997).
- [13] G. D. Hawkins, C. J. Cramer, and D. G. Truhlar, *Pairwise solute descreening of solute charges from a dielectric medium*, Chem. Phys. Lett. **246**, 122–129 (1995).
- [14] G. D. Hawkins, C. J. Cramer, and D. G. Truhlar, *Parametrized models of aqueous free energies of solvation based on pairwise descreening of solute atomic charges from a dielectric medium*, J. Phys. Chem. **100**, 19824–19839 (1996).
- [15] V. Tsui and D. A. Case, *Molecular dynamics simulations of nucleic acids with a generalized Born solvation model*, J. Am. Chem. Soc. **122**, 2489–2498 (2000).
- [16] V. Tsui and D. A. Case, *Theory and applications of the generalized Born solvation model in macromolecular simulations*, Biopolymers (Nucl. Acid Sci.) **56**, 275–291 (2001).
- [17] A. Onufriev, D. Bashford, and D. A. Case, *Exploring protein native states and large-scale conformational changes with a modified generalized Born model*, Proteins: Struct. Func. Bioinf. **55**, 383–394 (2004).
- [18] J. Srinivasan, M. W. Trevathan, P. Beroza, and D. A. Case, *Application of a pairwise generalized Born model to proteins and nucleic acids: Inclusion of salt effects*, Theor. Chem. Acc. **101**, 426–434 (1999).
- [19] O. Guvench, J. Weiser, P. Shenkin, I. Kolossváry, and W. C. Still, *Application of the frozen atom approximation to the GB/SA continuum model for solvation free energy*, J. Comput. Chem. **23**, 214–221 (2002).
- [20] J. Weiser, P. S. Shenkin, and W. C. Still, *Approximate atomic surfaces from linear combinations of pairwise overlaps (LCPO)*, J. Comput. Chem. **20**, 217–230 (1999).
- [21] L. Fejes Tóth, *Regular Figures*, vol. 48 of *International Series of Monographs in Pure and Applied Mathematics* (Macmillan, 1964).
- [22] N. Gō and H. A. Scheraga, *Ring closure and local conformational deformations of chain molecules*, Macromolecules **3**, 178–187 (1970).
- [23] L. R. Dodd, T. D. Boone, and D. N. Theodorou, *A concerted rotation algorithm for atomistic Monte Carlo simulation of polymer melts and glasses*, Mol. Phys. **78**, 961–996 (1993).
- [24] M. Mezei, *Efficient Monte Carlo sampling for long molecular chains using local moves, tested on a solvated lipid bilayer*, J. Chem. Phys. **118**, 3874–3879 (2003).
- [25] N. Gō and H. A. Scheraga, *Analysis of the contribution of internal vibrations to the statistical weights of equilibrium conformations of macromolecules*, J. Chem. Phys. **51**, 4751–4767 (1969).
- [26] V. Katritch, M. Totrov, and R. Abagyan, *ICFF: A new method to incorporate implicit flexibility into an internal coordinate force field*, J. Comput. Chem. **24**, 254–265 (2003).
- [27] M. W. Schmidt, K. K. Baldrige, J. A. Boatz, *et al.*, *The general atomic and molecular electronic structure system*, J. Comput. Chem. **14**, 1347–1363 (1993).